

Minutes of X3T11 ad hoc HIPPI SWG, and HNF Technical Committee (TC)
Special Working Meeting to Address HIPPI-6400 Issues Only
Tuesday & Wednesday, April 9-10, 1996
Palm Springs, California

1. Opening remarks and introductions

The Chairman, Don Tolmie of Los Alamos National Laboratory, opened this special working meeting at 1 PM by thanking Jeff Stai and Western Digital for hosting this meeting. This group is constituted as both the HIPPI special working group (SWG) under X3T11, and the HIPPI Networking Forum (HNF) - Technical Committee (TC).

Don Tolmie apologized for failing to complete minutes for the interim HIPPI-6400 meeting in Mountain View, CA. The incomplete action items from that meeting will be added to these minutes. When asked about whether minutes were needed from the interim meetings, attendees stated that at least the major agreements should be documented.

2. Review / modify the draft agenda

The draft agenda was distributed by e-mail before the meeting, and hard copies were at the meeting. No additional items were added. Don tried to organize the agenda with the more basic, and easiest, items first, with the more contentious items for the end, so we can address as much as possible in our short meeting time.

3. Review HIPPI-6400-PH changes since Rev 0.15

HIPPI-6400-PH Rev 0.2 was available via ftp a few days before the meeting, and was also distributed at the meeting. The major changes since Rev 0.15 were listed in the front of the document. The page numbers in the following headings refer to Rev 0.2.

3.1 Figure 2 bit and byte nomenclature (page 4)

After some discussion, the bit and byte nomenclature in figure 2 was accepted. It was noted that the document was not consistent throughout, some places have the bits and bytes in the wrong order.

In 6.1, the text in the first two paragraphs was changed so that DB00 and c00 are transmitted first. Also, control bit assignments for the micro-packet header were swapped around, e.g., VC changed from c63-c62 to c00-c01. The relative position of the TYPE and ABORT bits were swapped; the VCR and CR

fields were moved ahead of TSEQ; and the LCRC and ECRC were swapped. Don also agreed to add a figure showing the layout of the control bits in 16-bit quantities.

3.2 Signal line bit assignments (pages 16-17)

It was decided to swap the order that bytes were transmitted as detailed in tables 7 and 8. Namely, in signal line D00, bit d00.0 will go first, followed by d00.1, d00.2, d00.3, etc. The control signals are correct as is.

3.3 Bit nomenclature for DC balance (page 18)

The descriptive text is wrong when giving the examples for signal line D00. For example, at the end of the first paragraph of 10.2, it should be d00.0, d00.1, d00.2, and d00.3. The same applies to the list in the middle of the left column. The names w,x,T,y,z seemed to aid understanding.

In 10.3, the first note about the FRAME signal marking the first bit was removed.

3.4 Signal waveforms (pages 19-20)

In figures 6 and 7, the bit numbers for signal line D00 need to be renumbered as noted above. The w,x,T,y,z at the right end of figure 7 is incorrect.

Stan Swirhun noted that the FRAME signal has a very long period in relation to the bit signals and wondered if this would affect the dynamic skew compensation. Since all of the signals undergo the same long period during training, this may not be a problem. Stan will study it some more. Greg Chesson said that SGI will also consider other patterns for the FRAME and training signals.

It was agreed to delete clause 10.4--its information should be included in the optical and electrical clauses, not here. Don Tolmie will flesh out 10.5 to describe the FRAME signal some more. Some editing was done to the text in clause 11 on dynamic skew compensation.

It was agreed that dynamic skew compensation is mandatory, and a training sequence must occur at least every 10 μ s since it is also used to correct for

differences in clock rates--Dave Parry agreed to supply some explanatory text.

3.5 Rename "ABORT" to "ERROR" (page 8)

Don Tolmie gave some history behind the name ABORT, and stated that now the name ERROR seemed more appropriate. By unanimous agreement the name will be changed to ERROR throughout the document.

3.6 Review table 2 contents, (page 9)

The micro-packet contents summary was reviewed. "n.u.", meaning "not used, may contain any value", was felt to be difficult to test, and ambiguous. Many of these entries were changed to "0*", meaning transmit as 0's, but do not check at the receiver. Likewise, 1* means transmit as 1's but don't check, for example the TAIL bit, and FF* means transmit as FF but don't check.

Adding the Admin TYPE was accepted, and the details of the Admin data contents will be defined in the HIPPI-6400-SC document, or some other document--not HIPPI-6400-PH.

3.7 CRC operations (pages 10 and 15)

In 6.6.1, at the end of the first paragraph, Dave Parry suggested changing "micro-packet" to "data micro-packet" in two places. Wally St.John requested a global change to remove the hyphen in "micro-packet", and everyone agreed. The statement that the two 16-bit CRCs is equivalent to a 32-bit CRC will be changed to "stronger than a single 16-bit CRC".

It was agreed to continue using ECRC and LCRC rather than CRC1 and CRC2. Dave Parry noted that in some cases the ECRC was not an end-to-end checksum.

Greg Chesson presented some work that he had done on CRCs. He proposed CRC-16 ($x^{16} + x^{12} + x^5 + 1$) for the LCRC, and is considering ECRC = $x^{16} + x^{12} + x^3 + x + 1$. Greg has a program that generates all possible error combinations, and is comparing the results of the two CRCs to see if they share any common errors (which would be undetected). Greg proposed using the transmission order as the order that the ECRC bits would be entered into the calculation. Ed Grivna noted that the proper name for CRC-16 is ITU V.41. It was also pointed out that our 4b/5b coding may result in some error multiplication, e.g., a single bit line error resulting in 4 data bits in error. Greg will

continue to model the CRC behavior before proposing a final polynomial for the ECRC.

The question of whether or not the ECRC should accumulate over the header micropacket was discussed. It was originally proposed not to accumulate since the header may change due to source routing. It was now felt that source routing was not needed and its functions could be achieved in other ways. Greg Chesson said that it would probably still be a good idea not to accumulate over the header since the contents of the header micropacket and data micropackets would probably come from different places; no disagreement voiced.

A suggestion to capitalize "message" and "micro-packet" as proper names was met with opposition from those who felt that capital letters made the sentences harder to parse. It was agreed to capitalize if the word is used both in a general sense, and as a proper name; neither proposed word met this criteria.

At the end of the second paragraph of 6.6.1 the words about carrying the ECRC across intermediate devices will be expanded to better define what we determine an intermediate device to be.

6.6.2 will be changed so that the LCRC covers the ECRC.

In the last paragraph of 6.6.3, checking the ECRC prior to transmission, will be changed from optional to mandatory.

The pseudo-code in 9.1 was reviewed and modified. The action to "log error" was removed as being outside the scope of the standard. It was noted that the text in 9.1 describing retransmission was incorrect. There are a bunch of other errors that should be checked for as well (e.g., skipped TSEQ number, timeout, undefined TYPE value, VCR overrun) -- Bob Newhall agreed to supply some text describing them. 9.2 needs to be reworded.

3.8 Header micro-packet contents (pages 11-12)

Don Tolmie had proposed some fields (Schedule_Header and H800_Header) and a layout, based on his previous HIPPI-800 mapping proposal. Greg Chesson described some work by himself and Jim Pinkerton of SGI on a header for a message passing interface (MPI). Greg felt the need to keep the basic MPI header, and noted many similarities between it and Don's proposal. It was agreed to defer detailed discussion on the header micropacket

contents until Greg and Jim can try merging the two proposals.

It was agreed to use the words "chunk" and "microchunk" as the transfer units.

3.9 VC3 operations (pages 12-14)

It was agreed to defer work on this until the header was defined.

8.4 defining the VC priorities was removed from the document based on the reality that you cannot dictate a Source's behavior. This is not the case for a switch, where something will need to be specified.

3.10 Proposals for Admin micro-packet functions

No specific proposals were presented. When they are, the information should go into the HIPPI-6400-SC document, or at least a document other than HIPPI-6400-PH.

3.11 Reset, Initialize, Re-training, Warm Start, Cold Start (pages 8-9 and page 21)

Greg Chesson presented a HIPPI-6400 Startup Procedure detailing the sequences used for the different operations. The group reviewed the proposal and made some suggestions for changes. A state diagram will be added to help explain the procedures.

The Reset and Initialize sequences use essentially the same procedure; the difference is that the Initialize may propagate to other nodes while Reset just affects the local link. Initialize is what had originally been called "Cold start", and Reset is what had been called "Warm start". Bob Newhall agreed to provide some text for what gets reset, and when, for each.

4. HIPPI-6400-SC

Roger Ronald provided a first draft of HIPPI-6400-SC and a partial read-through was done. It was agreed that we would not support mandatory or optional source routing. Alternate paths will not be supported. Editorial work needs to be done to make sure that requirements use the word "shall", and things that are specified in HIPPI-6400-PH are not also specified in HIPPI-6400-SC.

Rather than interleaving every 16 micropackets, it will be specified as occurring at least every 65 micropackets. In-band switch management will be

required. The ECRC will not be recalculated in the fabric.

5. Discussion of copper interface (7:30 PM)

This discussion was a kick-off to specifying the copper interface for HIPPI-6400-PH. The folks in the FC-0 Copper meeting were specifically invited, and many attended.

5.1 Brief overview of requirements

Don Tolmie passed out a brief description of the HIPPI-6400-PH copper details as presently known. For example, the width was proposed as 16+4+1+1 for copper, and 8+2+1+1 for fiber. This was based on a lower frequency for the copper interface. Waveforms were also included.

5.2 Discussion of what is possible

Ed Grivna said that the narrower interface may be better since the cable would not be as bulky. Ed also suggested 150 ohm twisted-pair or 75 ohm coax. He said that EMC may be a problem. Stan Swirhun felt that the frequency difference between the FRAME signal and data bits may be a problem. Stan will study it further. Greg Chesson said that SGI would look at other encodings for FRAME.

5.3 Work planning

Henry Brandt agreed to lead a group in defining the copper interface. He will contact potential vendors before the Santa Fe meeting and see what they feel they have available for our needs. A meeting to address the topic, and get presentations from potential vendors, is scheduled for Tuesday, June 11, 1996, 6 PM - 9 PM, in Santa Fe, NM. The present HIPPI-6400-PH experts were requested to attend to input requirements, but it was acknowledged that they were not the copper experts. Don Tolmie agreed to put the copper guidelines on e-mail.

6. Planning for work on the optical interface

Tim Clay of Methode Electronics described a connector for parallel fiber that fits in an SC connector footprint. It is presently a prototype.

Some fiber vendors also requested presentation time in Santa Fe, and it was agreed to also let them have some time on Tuesday before the copper presentations. It was felt that copper was needed

HIPPI-6400 SWG meeting minutes - April 9-10, 1996

immediately for HIPPI-6400-PH, and fiber could come later.

Stan Swirhun and Jonathan Thatcher agreed to co-lead the group developing the fiber interface specifications.

7. Other "Open Issues" that haven't been covered yet

No other open issues were covered due to lack of time at the meeting. The schedule that was used seemed to work OK. Namely, Monday: HIPPI-6400 1 PM - 7 PM with pizza brought in for supper (thanks Steve Quan), 7 PM - 9 PM copper and some fiber; Tuesday: other HIPPI and HNF plenary 8 AM - noon, HIPPI-6400 1:30 PM - 5 PM. We had the room available till 9 PM but quit due to being pooped.

8. Future meeting schedule

May 9-10, 1996, Dallas, TX

Thursday, May 9 -

1 PM - 9 PM — HIPPI-6400

Friday, May 10 -

8AM - 2 PM — HIPPI-6400

The location is the Hyatt Regency DFW in Dallas, TX. The hotel is at the DFW airport so rental cars should not be needed. The phone number is (214) 615-6891, Fax (214) 615-6829, and the room rate is \$119 + 12% tax. The group name when making reservations is "E-Systems", and the cutoff date is April 25. Roger Ronald and E-Systems are the host.

June 10-11, 1996, Santa Fe, NM

Monday, June 10 -

1 PM - 9 PM — HIPPI-6400

Tuesday, June 11 -

8 AM - 3 PM — HIPPI General , HNF, HIPPI-6400

3 PM - 6 PM — HIPPI-6400 Optical

6 PM - 9 PM — HIPPI-6400 Copper

The location is the La Fonda Hotel, 100 E. San Francisco, Santa Fe, NM 87501, phone (505) 982-5511, Fax (505) 988-2952. The rate is \$105 single or double, parking and tax included. The closing data for reservations is May 10. The group name is X3T11-Los Alamos. Don Tolmie and Los Alamos are the host.

The schedule for the rest of 1996 is listed below. The Plenary meetings include HIPPI-6400, an HNF plenary, and all other HIPPI items. The Interim meeting cover just HIPPI-6400 items.

Jul 10-11	Interim	Mountain View, CA	SGI
Aug 5-6	Plenary	Honolulu, HI	Hitachi
Sep 11-12	Interim	Albuquerque, NM	Los Alamos
Oct 7-8	Plenary	St. Petersburg Beach, FL	AMP
Nov 6-7	Interim	Phoenix, AZ	Loral
Dec 2-3	Plenary	Rochester, MN	IBM

9. Review action items

1. Marti Bancroft to draft an annex describing problems and potentials for different VC priority schemes.
2. Stan Swirhun, and others, to consider problems with FRAME signal frequency.
3. Greg Chesson and others at SGI to consider other codings for the FRAME signal to avoid the long steady periods.
4. Dave Parry to provide some explanatory text about the 10 μ s requirement for training sequences.
5. Greg Chesson to check potential CRC codes for goodness, and propose one for the ECRC.
6. Bob Newhall to provide list of errors that will be detected by SuMAC.
7. Bob Newhall to provide descriptions of what gets reset by Reset and Initialize, and when.
8. Greg Chesson and Jim Pinkerton to merge the MPI and Schedule headers and propose a general micropacket header.
9. Henry Brandt to contact potential copper vendors before the Santa Fe meeting, explain the problem, and solicit their participation in specifying the copper interface.
10. Don Tolmie to make copper guidelines available via e-mail.
11. Roger Ronald to update HIPPI-6400-SC Rev 0.01 with changes agreed to at this meeting.
12. Don Tolmie to update HIPPI-6400-PH Rev 0.2 with changes agreed to at this meeting.

10. Adjournment

The meeting adjourned at 5 PM on the second day after a long, intense, and fruitful meeting.

HIPPI-6400 SWG meeting minutes - April 9-10, 1996

Attendance

Michael Griffin	3M Co.
Scott Stevens	ComScope, General Instrument
Tim Johnson	Cray Research
Ed Grivna	Cypress Semiconductor
Michael McGowen	Essential Communications
Roger Ronald	E-Systems
Francois Gaullier	Hewlett-Packard
Steve Joiner	Hewlett-Packard OCD
Ken Hahn	Hewlett-Packard
Christie Rice	Honeywell
Henry Brandt	IBM
Jonathan Thatcher	IBM
Barry Wightman	Lion Rock, Inc.
Chris Olson	Loral Defense Systems
James Hoffman	Los Alamos National Lab
Wally St. John	Los Alamos National Lab
Don Tolmie	Los Alamos National Lab
Tim Clay	Methode Electronics
Pat Gilliland	Methode Electronics
Richard O'Connell	Myriad Logic
Stephen Quan	NASA Ames Research Center
John Renwick	NetStar
Joe Parker	Optivision
Clive Towndrow	PsiTech Inc.
Art Beckman	Silicon Graphics
Greg Chesson	Silicon Graphics
Bob Newhall	Silicon Graphics
Dave Parry	Silicon Graphics
James Pinkerton	Silicon Graphics
Scott Kurimoto	Sumisho Electronics
Isao Morishita	Sumisho Electronics
Brian Hackler	Thomas & Betts
Don Knasel	US Connector
Stan Swirhun	Vixel Corp.